arXiv:1612.01527v2 [cs.CC] 12 Dec 2016

Matrix multiplication algorithms from group orbits

Joshua A. Grochow*† and Cristopher Moore*

December 13, 2016

Abstract

We show how to construct highly symmetric algorithms for matrix multiplication. In particular, we consider algorithms which decompose the matrix multiplication tensor into a sum of rank-1 tensors, where the decomposition itself consists of orbits under some finite group action. We show how to use the representation theory of the corresponding group to derive simple constraints on the decomposition, which we solve by hand for $n = 2, 3, 4, 5$, recovering Strassen's algorithm (in a particularly symmetric form) and new algorithms for larger n. While these new algorithms do not improve the known upper bounds on tensor rank or the matrix multiplication exponent, they are beautiful in their own right, and we point out modifications of this idea that could plausibly lead to further improvements. Our constructions also suggest further patterns that could be mined for new algorithms, including a tantalizing connection with lattices. In particular, using lattices we give the most transparent proof to date of Strassen's algorithm; the same proof works for all n, to yield a decomposition with $n^3 - n + 1$ terms.

1 Introduction

The complexity of matrix multiplication is a central question in computational complexity, bearing on the complexity not only of most problems in linear algebra, but also of myriad combinatorial problems, e.g., various shortest path problems [Zwi02] and bipartite matching problems [San09]. The main question around matrix multiplication is whether two $n \times n$ matrices can be multiplied in time $O(n^{2+\varepsilon})$ for every $\varepsilon > 0$; the current best upper bound on this exponent is 2.3728639 [LG14], narrowly beating [DS13, Wil12]. The best known lower bound is still only $3n^2 - o(n)$ [Lan14].

Since Strassen's 1969 paper [Str69], which showed how to beat the standard $O(n^3)$-time algorithm, it has been understood that one way to get asymptotic improvements in algorithms for matrix multiplication is to find algebraic algorithms for multiplying small, fixed-size matrices using few multiplications, and then to apply these algorithms recursively. In this note, we are thus primarily interested in algebraic algorithms for matrix multiplication.

Recently, there have been several papers analyzing the geometry and symmetries of algebraic algorithms for small matrices [Bur14, Bur15, LR16, LM16b, CILO16]. Here, we take this line of research one step further by using symmetries to discover *new* algorithms for multiplying matrices of small size. While our algorithms do not improve the state-of-the-art bounds on the matrix multiplication exponent, they suggest that we can use group symmetries and group orbits to find new algorithms for matrix multiplication. In addition to their potential value for future endeavors, we believe that these highly symmetric matrix multiplication algorithms are beautiful in their own right, and deserve to be shared simply for their beauty.

Although this line of research was also suggested in [CILO16], who analyze the symmetries of Strassen's algorithm and promise a forthcoming paper with a new construction along these lines for $n = 3$, our paper is the result of an independent project started two years ago, and we find constructions for $n = 3$, 4, and 5. Moreover, these constructions appear to yield a larger pattern. In the final section, we show how—although

*Santa Fe Institute, 1399 Hyde Park Road, Santa Fe, New Mexico 87501 USA

†Department of Computer Science, University of Colorado at Boulder, 1111 Engineering Drive, ECOT 717, 430 UCB, Boulder, Colorado 80309 USA

our *method* of construction is more general than this—the constructions we ended up finding were in fact all instances of a single lattice-based construction yielding rank $n^3 - n + 1$ for $n \times n$ matrix multiplication. The proof that the lattice-based construction works is perhaps the simplest and most transparent proof of Strassen's algorithm that we are aware of, and the same proof works for all n.

2 Background

For general background on algebraic complexity, we refer the reader to the book [BCS97]. Bläser's survey article [Blä13], in addition to excellent coverage around matrix multiplication, has a nice tutorial on tensors and their basic properties.

In the algebraic setting, since matrix multiplication is a bilinear map, it is known that it can equivalently be reformulated as a tensor, and that the algebraic complexity of matrix multiplication is within a factor of 2 of the rank of this tensor. The matrix multiplication tensor for $n \times n$ matrices is

$$\mathrm{MM}^{abc}_{def} = \delta^a_e \delta^b_f \delta^c_d \,,$$

where the indices range from 1 to n. This is also defined by the inner product

$$\langle \mathrm{MM} \,|\, A \otimes B \otimes C \rangle = \operatorname{tr} ABC \,.$$

Given vector spaces $V_1, \ldots, V_k$, a vector $v \in V_1 \otimes \cdots \otimes V_k$ is said to have tensor rank one if it is a separable tensor, that is, of the form $v_1 \otimes v_2 \otimes \cdots \otimes v_k$ for some $v_i \in V_i$. The *tensor rank* of v is the smallest number of rank-one tensors whose sum is v. In the case of MM, we have $k = 3$ and $V_1 \cong V_2 \cong V_3 \cong \mathbb{C}^{n \times n}$.

Although it won't be relevant for our results, we cannot resist mentioning border rank. The border rank of a tensor $v \in V_1 \otimes \cdots \otimes V_k$ is the smallest r such that v is a limit of tensors of rank at most r. The border rank is less than or equal to the rank, and indeed it can be strictly smaller. However, the exponent of matrix multiplication is the same whether measured by tensor rank or border rank [Bin80].

Here is a brief summary of what is known about the tensor and border rank of small (square) matrix multiplication tensors:

n	Tensor rank $\geq$	Tensor rank $\leq$	Border rank $\geq$	Border rank $\leq$
2	7 [HK71, Win71]	7 [Str69]	7 [Lan06, HIL13]	7 [Str69]
3	19 [Blä03]	23 [Lad76, JM86, CBH11]	16 [LM16b]	21 [Sch81]
4	34 [Blä03]	49 [Str69]	29 [LM16b]	49 [Str69]

The tensor rank of $n \times n$ matrix multiplication is at least $3n^2 - o(n)$ [Lan14]. The previous best bound was $(5/2)n^2 - 3n$ [Blä99], and the $3n^2 - o(n)$ beats this out for $n > 84$. The border rank of $n \times n$ matrix multiplication is at least $2n^2 - \lceil \log_2 n \rceil - 1$ [LM16a] (and at least $2n^2 - n + 1$ [LM16b], which is a tighter lower bound for all $n < 6$).

The current world record algorithm [LG14] and its immediate precursors [DS13, Wil12] were achieved by analyzing powers of the Coppersmith–Winograd tensor [CW90]. Ambainis, Filmus, and Le Gall [AFLG15] showed that taking further powers of the Coppersmith–Winograd tensor cannot get an exponent better than 2.3725, and that even certain generalized versions of this method cannot beat exponent 2.3078. The smallest square cases which could beat these bounds (world record, 2.3725, and 2.3078, respectively), based on the lower bounds above, are:

- Showing the border rank of 6×6 matrix multiplication is at most 70 would beat the world record (and even exponent 2.3725). This must be border rank, as the above bounds show that its rank is at least 72. (Its border rank is at least 68.)
- Showing the tensor rank of 8×8 matrix multiplication is at most 138 would beat the world record (and 2.3725). Its rank is at least 136, and the current best upper bound is $7^3 = 343$ (the tensor cube of Strassen's algorithm).

- Showing the border rank of 9×9 matrix multiplication is at most 159 would beat the stronger bound of 2.3078. This must be border rank, as its rank is at least 176. (Its border rank is at least 157.)
- Showing that the tensor rank of 16×16 matrix multiplication is at most 600 would beat the stronger bound of 2.3078. Its rank is at least 592, and the current best upper bound is $7^4 = 2401$ (the fourth tensor power of Strassen's algorithm).

3 Using Fourier analysis to rediscover Strassen as a group orbit

We begin by "Fourier analyzing" the matrix multiplication tensor. First recall that MM is defined by its symmetries (see, e.g., [BI11]), that is, up to a constant, it is the unique operator fixed under the following action of $\mathsf{GL}(n)^3$. For, given $X, Y, Z \in \mathsf{GL}(n)$, we have

$$\mathrm{MM} = (X \otimes Y \otimes Z)\mathrm{MM}(Z^{-1} \otimes X^{-1} \otimes Y^{-1})\,, \tag{1}$$

where hopefully the notation is clear. To see that MM has this symmetry, note that

$$\operatorname{tr} ABC = \operatorname{tr}(Z^{-1}AX)(X^{-1}BY)(Y^{-1}CZ)\,.$$

To see that MM is the only such operator up to a constant, we simply note that if $\rho \cong \mathbb{C}^n$ is the standard representation of $\mathsf{GL}(n)$, then $\rho^{\otimes 3}$ is an irreducible representation of $\mathsf{GL}(n)^3$. Taking the tensor product with its dual $(\rho^{\otimes 3})^*$ then gives a single copy of the trivial representation, so MM lives in—and therefore spans—this one-dimensional subspace.

The same argument applies to any group G that has an n-dimensional irreducible presentation, which we will also call ρ. Since $\rho^{\otimes 3}$ is an irrep of G^3, $\rho^{\otimes 3} \otimes (\rho^{\otimes 3})^*$ contains a unique one-dimensional subspace of operators invariant under the action of G^3 where X, Y, Z are each $\rho(g)$ for some $g \in G$. This suggests that a good way to search for new matrix multiplication algorithms is to start with sums of separable tensors where the sum has some symmetry built in from the beginning.

We can also consider the action of the diagonal subgroup $G = \{(g, g, g) : g \in G\} \subset G^3$, which conjugates each factor in the tensor product by the same group element. Clearly MM is invariant under this action, but it is no longer the only such operator, since $\rho^{\otimes 3}$ now has multiple copies of the trivial representation when acted on by this smaller group. The idea, independently suggested in [CILO16], is to study how MM lies across these various trivial subspaces. In this section we carry out this program in the case $n = 2$ and $G = S_3$, where ρ is the standard irrep of S_3.

The conjugation action of S_3 on the space of 2×2 matrices is $\rho \otimes \rho^* \cong \mathbf{1} \oplus \pi \oplus \rho$ where $\mathbf{1}$ and π are the trivial and parity representations respectively. To make things more concrete, we choose a basis for ρ, which corresponds to the rotations and reflections of an equilateral triangle centered on the origin, whose corners labeled $1, 2, 3$ lie at $(1, 0)$, $(-1/2, \sqrt{3}/2)$, and $(-1/2, -\sqrt{3}/2)$. Then the generators of the group correspond to the unitary matrices

$$\rho((23)) = \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix} \triangleq \phi \quad \text{and} \quad \rho((123)) = \begin{pmatrix} -1/2 & -\sqrt{3}/2 \\ \sqrt{3}/2 & -1/2 \end{pmatrix} \triangleq \sigma\,, \tag{2}$$

where we think of these as acting on column vectors on the left. Geometrically, ϕ flips the plane around the x-axis, and σ rotates it $2\pi/3$ counterclockwise.

The subspaces corresponding to $\mathbf{1}$, π, and ρ are spanned by the following vectors in $\rho \otimes \rho^*$, which we denote by $\dim \rho \times \dim \rho^* = 2 \times 2$ matrices:

$$\mathbf{1} = \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix} \quad \pi = \begin{pmatrix} 0 & -1 \\ 1 & 0 \end{pmatrix} \quad \rho_x = \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix} \quad \rho_y = \begin{pmatrix} 0 & -1 \\ -1 & 0 \end{pmatrix},$$

where we abuse notation by identifying the irreps $\mathbf{1}$ and π with matrices that span them. The reader can check that

$$\phi\pi\phi^\dagger = -\pi \quad \text{and} \quad \sigma\pi\sigma^{-1} = \pi\,, \tag{3}$$

confiming that π spans the parity representation. We have chosen ρ_x and ρ_y so that they correspond to the basis vectors $(1,0)$ and $(0,1)$ in the original basis for ρ. In particular,

$$\phi\rho_x\phi^\dagger = \rho_x\,, \quad \phi\rho_y\phi^\dagger = -\rho_y\,, \quad \sigma\rho_x\sigma^\dagger = -\frac{1}{2}\rho_x + \frac{\sqrt{3}}{2}\rho_y\,, \quad \text{and} \quad \sigma\rho_y\sigma^\dagger = -\frac{\sqrt{3}}{2}\rho_x - \frac{1}{2}\rho_y\,, \tag{4}$$

(We know that $\rho_x = \phi$, but we want to distinguish ρ_x, which is a vector in $\rho \otimes \rho^*$, from ϕ, which is an operator acting on that space by conjugation.)

We now have a convenient basis $\{\mathbf{1}, \pi, \rho_x, \rho_y\}^{\otimes 3}$ for the space of triples of 2×2 matrices, i.e., for $(\rho\otimes\rho^*)^{\otimes 3}$. Next we write MM in this basis. Since these 64 basis vectors are orthogonal (under the unique S_3-invariant inner product of $A, B \in \rho \otimes \rho^*$ defined as $\langle A\,|\,B\rangle = \operatorname{tr}(A^\dagger B)$), we have

$$\mathrm{MM} = \sum_{(\alpha,\beta,\gamma)\in\{\mathbf{1},\pi,\rho_x,\rho_y\}^3} c_{\alpha,\beta,\gamma}\,\alpha\otimes\beta\otimes\gamma\,,$$

where

$$c_{\alpha,\beta,\gamma} = \frac{\langle \mathrm{MM}\,|\,\alpha\otimes\beta\otimes\gamma\rangle}{\langle \alpha\otimes\beta\otimes\gamma\,|\,\alpha\otimes\beta\otimes\gamma\rangle} = \frac{1}{8}\operatorname{tr}\alpha\beta\gamma\,.$$

(The 1/8 here comes from the fact that $\langle\alpha\otimes\beta\otimes\gamma\,|\,\alpha\otimes\beta\otimes\gamma\rangle = \langle\alpha\,|\,\alpha\rangle\,\langle\beta\,|\,\beta\rangle\,\langle\gamma\,|\,\gamma\rangle$, and for all $\alpha \in \{\mathbf{1}, \pi, \rho_x, \rho_y\}$ we have $\langle\alpha\,|\,\alpha\rangle = 2$.) This gives

$$\begin{aligned}
\mathrm{MM} = \frac{1}{4}\big(&\mathbf{1}\otimes\mathbf{1}\otimes\mathbf{1} \\
&- \mathbf{1}\otimes\pi\otimes\pi - \pi\otimes\mathbf{1}\otimes\pi - \pi\otimes\pi\otimes\mathbf{1} \\
&+ \mathbf{1}\otimes\rho_x\otimes\rho_x + \rho_x\otimes\mathbf{1}\otimes\rho_x + \rho_x\otimes\rho_x\otimes\mathbf{1} \\
&+ \mathbf{1}\otimes\rho_y\otimes\rho_y + \rho_y\otimes\mathbf{1}\otimes\rho_y + \rho_y\otimes\rho_y\otimes\mathbf{1} \\
&+ \pi\otimes\rho_x\otimes\rho_y + \rho_y\otimes\pi\otimes\rho_x + \rho_x\otimes\rho_y\otimes\pi \\
&- \pi\otimes\rho_y\otimes\rho_x - \rho_x\otimes\pi\otimes\rho_y - \rho_y\otimes\rho_x\otimes\pi\big)\,.
\end{aligned} \tag{5, 6}$$

We streamlined this calculation by noting that each of the 27 tensor products of representations $\{\mathbf{1}, \pi, \rho\}^{\otimes 3}$ has at most one copy of the trivial representation of the diagonal S_3 action, and that it is only these triples for which we can have $c_{\alpha,\beta,\gamma} \neq 0$. In particular, the third and fourth lines correspond to the trivial subspace of $\mathbf{1}\otimes\rho\otimes\rho$, which is spanned by

$$\mathbf{1}\otimes\psi_{\mathbf{1}}^{\rho\rho} \quad \text{where} \quad \psi_{\mathbf{1}}^{\rho\rho} = \rho_x\otimes\rho_x + \rho_y\otimes\rho_y\,,$$

since $\psi_{\mathbf{1}}^{\rho\rho}$ spans the trivial subspace of $\rho\otimes\rho$. Similarly, the fifth and sixth lines correspond to the trivial subspace of $\pi\otimes\rho\otimes\rho$, which is spanned by

$$\pi\otimes\psi_{\pi}^{\rho\rho} \quad \text{where} \quad \psi_{\pi}^{\rho\rho} = \rho_x\otimes\rho_y - \rho_y\otimes\rho_x\,,$$

since $\psi_{\pi}^{\rho\rho}$ spans the parity subspace of $\rho\otimes\rho$. Interestingly, $\rho\otimes\rho\otimes\rho$ does contain a copy of the trivial representation, spanned by

$$\psi_{\mathbf{1}}^{\rho\rho\rho} = \rho_x\otimes\rho_x\otimes\rho_x - \rho_x\otimes\rho_y\otimes\rho_y - \rho_y\otimes\rho_x\otimes\rho_y - \rho_y\otimes\rho_y\otimes\rho_x = \frac{(\rho_x + i\rho_y)^{\otimes 3} + (\rho_x - i\rho_y)^{\otimes 3}}{2}\,. \tag{7}$$

However, the corresponding coefficients in MM are all zero, since the product of any three $\rho \in \{\rho_x, \rho_y\}$ has zero trace.

While some of its terms can be combined, taking (6) directly does not lead to a low tensor rank for MM. We now set about retroactively discovering Strassen's algorithm, with the following ansatz:

$$\begin{aligned}
\mathrm{MM} &= \mathbf{1}\otimes\mathbf{1}\otimes\mathbf{1} + \sum_{g\in S_3} (\rho(g)m_1\rho(g^{-1}))\otimes(\rho(g)m_2\rho(g^{-1}))\otimes(\rho(g)m_3\rho(g^{-1})) \\
&= \mathbf{1}\otimes\mathbf{1}\otimes\mathbf{1} + \sum_{g\in S_3} \rho(g)^{\otimes 3}(m_1\otimes m_2\otimes m_3)\rho(g)^{\dagger\otimes 3}\,.
\end{aligned} \tag{8}$$

That is, besides the term $\mathbf{1}^{\otimes 3}$ which we generously grant ourselves, we assume that the other six terms consist of the orbit of some rank-1 tensor $m_1 \otimes m_2 \otimes m_3$ under the diagonal conjugation action of S_3. Moreover, we assume that m_1, m_2, m_3 form an orbit under the order-3 element $(1,2,3)$. Since $\sigma = \rho((123))$, this gives

$$m_1 = m\,, \quad m_2 = \sigma m \sigma^\dagger\,, \quad m_3 = \sigma^2 m (\sigma^\dagger)^2\,. \tag{9}$$

Note that this ansatz is manifestly invariant under both the diagonal conjugation action and the $\mathbb{Z}_3$ symmetry which rotates the three factors of the tensor product.

Now, since summing over an orbit annihilates any non-trivial irreps of S_3, it suffices for $m_1 \otimes m_2 \otimes m_3$ to have the correct projection into each copy of the trivial representation, which we enumerated above. This gives the following necessary and sufficient conditions:

$$\begin{aligned}
6\,\frac{\langle m_1 \otimes m_2 \otimes m_3 \,|\, \mathbf{1}\otimes\mathbf{1}\otimes\mathbf{1}\rangle}{\langle \mathbf{1}\otimes\mathbf{1}\otimes\mathbf{1} \,|\, \mathbf{1}\otimes\mathbf{1}\otimes\mathbf{1}\rangle} &= \frac{1}{4} - 1 \\
6\,\frac{\langle m_1 \otimes m_2 \otimes m_3 \,|\, \mathbf{1}\otimes\pi\otimes\pi\rangle}{\langle \mathbf{1}\otimes\pi\otimes\pi \,|\, \mathbf{1}\otimes\pi\otimes\pi\rangle} &= -\frac{1}{4} \\
6\,\frac{\langle m_1 \otimes m_2 \otimes m_3 \,|\, \mathbf{1}\otimes\psi_{\mathbf{1}}^{\rho\rho}\rangle}{\langle \mathbf{1}\otimes\psi_{\mathbf{1}}^{\rho\rho} \,|\, \mathbf{1}\otimes\psi_{\mathbf{1}}^{\rho\rho}\rangle} &= \frac{1}{4} \\
6\,\frac{\langle m_1 \otimes m_2 \otimes m_3 \,|\, \pi\otimes\psi_{\pi}^{\rho\rho}\rangle}{\langle \pi\otimes\psi_{\pi}^{\rho\rho} \,|\, \pi\otimes\psi_{\pi}^{\rho\rho}\rangle} &= \frac{1}{4} \\
6\,\frac{\langle m_1 \otimes m_2 \otimes m_3 \,|\, \psi_{\mathbf{1}}^{\rho\rho\rho}\rangle}{\langle \psi_{\mathbf{1}}^{\rho\rho\rho} \,|\, \psi_{\mathbf{1}}^{\rho\rho\rho}\rangle} &= 0\,,
\end{aligned}$$

or more concretely (where we use $\pi^\dagger = -\pi$, $\rho_x^\dagger = \rho_x$, and $\rho_y^\dagger = \rho_y$)

$$\begin{aligned}
(\operatorname{tr} m_1)(\operatorname{tr} m_2)(\operatorname{tr} m_3) &= -1 \\
(\operatorname{tr} m_1)(\operatorname{tr} \pi m_2)(\operatorname{tr} \pi m_3) &= -\frac{1}{3} \\
(\operatorname{tr} m_1)\big((\operatorname{tr} \rho_x m_2)(\operatorname{tr} \rho_x m_3) + (\operatorname{tr} \rho_y m_2)(\operatorname{tr} \rho_y m_3)\big) &= \frac{2}{3} \\
(\operatorname{tr} \pi m_1)\big((\operatorname{tr} \rho_x m_2)(\operatorname{tr} \rho_y m_3) - (\operatorname{tr} \rho_y m_2)(\operatorname{tr} \rho_x m_3)\big) &= -\frac{2}{3} \\
\big(\operatorname{tr}(\rho_x + i\rho_y) m_1\big)\big(\operatorname{tr}(\rho_x + i\rho_y) m_2\big)\big(\operatorname{tr}(\rho_x + i\rho_y) m_3\big) + \big(\operatorname{tr}(\rho_x - i\rho_y) m_1\big)\big(\operatorname{tr}(\rho_x - i\rho_y) m_2\big)\big(\operatorname{tr}(\rho_x - i\rho_y) m_3\big) &= 0\,.
\end{aligned}$$

Note that the rotated versions of these equations hold automatically given our assumption (9), since these are trivial subspaces under the diagonal conjugation action and, e. g., $m_2 \otimes m_3 \otimes m_1 = \sigma^{\otimes 3}(m_1 \otimes m_2 \otimes m_3)\sigma^{\dagger\otimes 3}$. Moreover, using (3) and (4) and a few lines of algebra, the above equations simplify to

$$(\operatorname{tr} m)^3 = -1 \tag{10}$$

$$(\operatorname{tr} m)(\operatorname{tr} \pi m)^2 = -\frac{1}{3} \tag{11}$$

$$(\operatorname{tr} m) \times -\frac{1}{2}\big((\operatorname{tr} \rho_x m)^2 + (\operatorname{tr} \rho_y m)^2\big) = \frac{2}{3} \tag{12}$$

$$(\operatorname{tr} \pi m) \times -\frac{\sqrt{3}}{2}\big((\operatorname{tr} \rho_x m)^2 + (\operatorname{tr} \rho_y m)^2\big) = -\frac{2}{3} \tag{13}$$

$$(\operatorname{tr} \rho_x m)^3 - 3(\operatorname{tr} \rho_x m)(\operatorname{tr} \rho_y m)^2 = 0\,. \tag{14}$$

Since we can multiply m by any cube root of 1 without changing $m_1 \otimes m_2 \otimes m_3$, we are free to solve (10) by setting

$$\operatorname{tr} m = -1\,. \tag{15}$$

Then (12) and (13) imply

$$\operatorname{tr} \pi m = \frac{1}{\sqrt{3}}\,, \tag{16}$$

which solves (11) and gives

$$(\operatorname{tr} \rho_x m)^2 + (\operatorname{tr} \rho_y m)^2 = \frac{4}{3}\,. \tag{17}$$

If we write

$$m = \begin{pmatrix} m_{11} & m_{12} \\ m_{21} & m_{22} \end{pmatrix}\,,$$

then

$$(\operatorname{tr} m)^2 + (\operatorname{tr} \pi m)^2 = (m_{11} + m_{22})^2 + (m_{21} - m_{12})^2 = |m|^2 + 2 \det m \tag{18}$$

where $|m|^2 = \operatorname{tr} m^\dagger m$ denotes the Frobenius norm, while

$$(\operatorname{tr} \rho_x m)^2 + (\operatorname{tr} \rho_y m)^2 = (m_{11} - m_{22})^2 + (m_{12} + m_{21})^2 = |m|^2 - 2 \det m\,. \tag{19}$$

But (15), (16), and (17) give

$$(\operatorname{tr} m)^2 + (\operatorname{tr} \pi m)^2 = (\operatorname{tr} \rho_x m)^2 + (\operatorname{tr} \rho_y m)^2 = \frac{4}{3}\,,$$

so we have

$$\det m = 0 \quad \text{and} \quad |m|^2 = \frac{4}{3}\,.$$

It follows that m has rank 1, and that it can be written $m = |u\rangle\langle v|$ for some vectors u, v such that

$$\operatorname{tr} m = \langle v \,|\, u \rangle = -1\,, \quad |m|^2 = |u|^2 |v|^2 = \frac{4}{3}\,, \quad \text{and} \quad \operatorname{tr} \pi m = \langle v \,|\, \pi u \rangle = \frac{1}{\sqrt{3}}\,.$$

Since rescaling u and v to cu and $c^{-1}v$ preserves $|u\rangle\langle v|$ for any constant c, we can assume that $|u| = 1$ and $|v| = 2/\sqrt{3}$. In addition, the first two of these conditions imply that the angle between u and v has cosine $\langle v \,|\, u \rangle / (|u||v|) = -\sqrt{3}/2$, and since π rotates the plane counterclockwise by $\pi/2$ the third condition implies that this angle is $5\pi/6$. Thus we can write

$$u = (\cos\theta, \sin\theta) \quad \text{and} \quad v = \frac{2}{\sqrt{3}} \left(\cos(\theta + 5\pi/6), \sin(\theta + 5\pi/6)\right) = \frac{2}{3}(\sigma u - u)\,. \tag{20}$$

We can show that θ belongs to a finite set using (14), which is the only one of our Fourier-based equations we have not used so far. Substituting $m = |u\rangle\langle v|$ yields

$$\langle v \,|\, \rho_x u \rangle^3 - 3 \langle v \,|\, \rho_x u \rangle \langle v \,|\, \rho_y u \rangle^2 = -\frac{8}{3\sqrt{3}} \sin 6\theta = 0\,, \tag{21}$$

so θ is a integer multiple of $\pi/6$, and any such multiple yields a decomposition of MM. In particular, setting $\theta = 0$ gives

$$u = (1, 0) \quad \text{and} \quad v = (-1, \sqrt{3})\,, \tag{22}$$

in which case

$$m = \begin{pmatrix} -1 & \sqrt{3} \\ 0 & 0 \end{pmatrix}\,.$$

Finally, we note that if we consider the following set of unit vectors which form an orbit under S_3,

$$w_1 = (1, 0)\,, \quad w_2 = \left(-\frac{1}{2}, \frac{\sqrt{3}}{2}\right)\,, \quad w_3 = \left(-\frac{1}{2}, -\frac{\sqrt{3}}{2}\right)\,,$$

then we can express u and v in (22) as

$$u = w_1 \quad \text{and} \quad v = \frac{2}{3}(w_2 - w_1)\,.$$

As we will see below, a similar construction works for the next few values of n.

4 Constructions for larger n

Strassen's algorithm, or rather its symmetric form derived above, suggests a family of matrix multiplication algorithms. Namely, suppose we have a group G with an irreducible representation ρ of dimension n (which we assume for simplicity is real) and which contains an order-3 element σ. Then we consider decompositions of the matrix multiplication tensor consisting of $\mathbf{1}^{\otimes 3}$ plus an orbit of some rank-1 matrix $m = |u\rangle\langle v|$,

$$\begin{aligned} \mathrm{MM} &= \mathbf{1}\otimes\mathbf{1}\otimes\mathbf{1} + \sum_{g\in G} \rho(g)^{\otimes 3}\left(m\otimes \sigma m\sigma^{-1}\otimes \sigma^2 m\sigma^{\dagger 2}\right)\rho(g)^{\dagger\otimes 3} \\ &= \mathbf{1}\otimes\mathbf{1}\otimes\mathbf{1} + \sum_{g\in G} \rho(g)^{\otimes 3}\left(|u\rangle\langle v|\otimes|\sigma u\rangle\langle\sigma v|\otimes|\sigma^2 u\rangle\langle\sigma^2 v|\right)\rho(g)^{\dagger\otimes 3} \end{aligned} \tag{23}$$

where u, v are vectors in $\mathbb{R}^n$ and $\sigma = \rho(h)$. Note that any decomposition of this form is manifestly symmetric under the diagonal conjugation action of G, as well as the $\mathbb{Z}_3$ action that rotates the three factors in the tensor product; both these actions simply permute the terms in the sum. The question is whether there are u, v such that (23) is also symmetric under the G^3 action, in which case it is a valid decomposition of the matrix multiplication tensor.

The constructions we have found have $|G| = n^3 - n$, giving a decomposition of tensor rank $n^3 - n + 1$. For $n > 2$ this is not optimal: it yields decompositions of rank 25 and 61 for $n = 3$ and $n = 4$ respectively, which exceed the known upper bounds of 23 [Lad76, JM86, CBH11] and 49 [Str69]. However, we hope that similar constructions involving multiple orbits will yield optimal decompositions in the future. For instance, the tensor square of Strassen's algorithm can be described in terms of the four-dimensional irrep of S_3^2, with one orbit of 36 rank-1 matrices, two orbits of 6 rank-2 matrices, and $\mathbf{1}^{\otimes 3}$, for a total tensor rank of 49.

We start by writing some necessary conditions on u and v. Viewed as a linear operator on $(\mathbb{R}^n)^{\otimes 3}$, the matrix multiplication tensor has trace $\operatorname{tr}\mathrm{MM} = n$ and Frobenius norm $|\mathrm{MM}|^2 = \langle \mathrm{MM}|\mathrm{MM}\rangle = n^3$. Finally, we have $\operatorname{tr}\mathrm{MM}^2 = \operatorname{tr}\mathrm{MM} = \left\langle \mathrm{MM}\middle|\mathrm{MM}^\dagger\right\rangle = n$. Since the trace is basis-invariant, the traces of the terms in (23) are

$$\begin{gathered} \operatorname{tr}\mathbf{1}^{\otimes 3} = n^3 \\ \operatorname{tr}\left(|u\rangle\langle v|\otimes|\sigma u\rangle\langle\sigma v|\otimes|\sigma^2 u\rangle\langle\sigma^2 v|\right) = \left(\operatorname{tr}|u\rangle\langle v|\right)^3 = \langle v|u\rangle^3 , \end{gathered}$$

and the inner product of each term with MM is

$$\begin{gathered} \left\langle \mathrm{MM}\middle|\mathbf{1}^{\otimes 3}\right\rangle = \operatorname{tr}\mathbf{1} = n \\ \left\langle \mathrm{MM}\middle||u\rangle\langle v|\otimes|\sigma u\rangle\langle\sigma v|\otimes|\sigma^2 u\rangle\langle\sigma^2 v|\right\rangle = \operatorname{tr}\left(|u\rangle\langle v\,|\sigma u\rangle\langle\sigma v\,|\sigma^2 u\rangle\langle\sigma^2 v|\right) = \langle v|\sigma u\rangle^3 , \end{gathered}$$

where we have used the fact that σ is unitary. Finally we have

$$\begin{gathered} \left\langle \mathrm{MM}\middle|\mathbf{1}^{\dagger\otimes 3}\right\rangle = \operatorname{tr}\mathbf{1} = n \\ \left\langle \mathrm{MM}\middle|\left(|u\rangle\langle v|\otimes|\sigma u\rangle\langle\sigma v|\otimes|\sigma^2 u\rangle\langle\sigma^2 v|\right)^\dagger\right\rangle = \operatorname{tr}\left(|v\rangle\langle u\,|\sigma v\rangle\langle\sigma u\,|\sigma^2 v\rangle\langle\sigma^2 u|\right) = \langle v|\sigma^2 u\rangle^3 . \end{gathered}$$

Thus (23) gives

$$n = n^3 + |G|\langle v|u\rangle^3 , \quad n^3 = n + |G|\langle v|\sigma u\rangle^3 , \quad \text{and} \quad n = n + \langle v|\sigma^2 u\rangle^3 . \tag{24}$$

As before, we use our freedom to multiply, say, u by a cube root of unity. Along with $|G| = n^3 - n$, this gives the conditions

$$\langle v|u\rangle = -1 , \quad \langle v|\sigma u\rangle = +1 , \quad \langle v|\sigma^2 u\rangle = 0 . \tag{25}$$

For instance, suppose $n = 2$, ρ is the standard representation of $G = S_3$, and u is a basis vector $(1, 0)$. If σ rotates $\mathbb{R}^2$ by $2\pi/3$ counterclockwise, the unique vector v that satisfies (25) is $(-1, 1/\sqrt{3})$, recovering the pair u, v we constructed for Strassen's algorithm (22) in the previous section.

More generally, suppose an order-3 element σ rotates a two-dimensional subspace R_σ by $2\pi/3$, and fixes its $(n-2)$-dimensional orthogonal complement $\overline{R_\sigma}$. We know from (25) that $u \neq \sigma u$, so u has a nonzero projection into R_σ. Since we only care about the product $m = |u\rangle\langle v|$, we can rescale u so that this projection is of unit length; denote it by y. We can then write $u = y + u_\perp$ where $u_\perp \in \overline{R_\sigma}$. Similarly write $v = z + v_\perp$ where $z \in R_\sigma$ and $v_\perp \in \overline{R_\sigma}$.

Since $u + \sigma u + \sigma^2 u = 3u_\perp$, summing the three conditions of (25) implies that $\langle v\,|\,u_\perp\rangle = \langle v_\perp\,|\,u_\perp\rangle = 0$. Therefore, $\langle v\,|\,u\rangle = \langle z\,|\,y\rangle$, $\langle v\,|\,\sigma u\rangle = \langle z\,|\,\sigma y\rangle$, and $\langle v\,|\,\sigma^2 u\rangle = \langle z\,|\,\sigma^2 y\rangle$. But since y is of unit length, the same argument that led to (20) shows that the unique vector $z \in R_\sigma$ such that $\langle z\,|\,y\rangle = -1$, $\langle z\,|\,\sigma y\rangle = +1$, and $\langle z\,|\,\sigma^2 y\rangle = 0$ is

$$z = \frac{2}{3}(\sigma y - y)\,. \tag{26}$$

Thus we have

$$u = y + u_\perp \quad \text{and} \quad v = z + v_\perp \tag{27}$$

where $y \in R_\sigma$ has unit length, $z \in R_\sigma$ is given by (26), and $\langle u_\perp\,|\,v_\perp\rangle = 0$.

This greatly reduces the dimensionality of our search, to choosing y and finding $u_\perp$ and $v_\perp$. We can reduce the search further by considering how $|u\rangle\langle v|$ projects into the various trivial subspaces of G. As we did for $n = 2$ above, it will be convenient to do this "Fourier analytically," that is, by breaking $\rho \otimes \rho^*$ up into irreps under the G^3 conjugation action and using them to define a judicious choice of basis for $(\rho \otimes \rho^*)^{\otimes 3}$. While for larger groups this calculation becomes rather lengthy, note that because of our ansatz (23), we only have $2n-1$ parameters to solve for. The conditions (25), resulting in the partial solution (27), reduce the number of parameters even further: there is a one-parameter family of unit vectors y in R_σ, and $u_\perp$ and $v_\perp$ together have $(n-2)+(n-3) = 2n-5$ free parameters. As a result it suffices to use just a few of the projections onto irreps to reduce to a finite set of possible solutions.

4.1 $n = 3$: the standard representation of S_4

For $n = 3$, we let ρ be the standard representation of $G = S_4$, which is of order $|S_4| = 24 = n^3 - n$. A convenient basis for this representation treats group elements as signed permutation matrices, whose determinants are the parities of the corresponding permutations. For instance,

$$\rho((12)) = \begin{pmatrix} 0 & 1 & 0 \\ 1 & 0 & 0 \\ 0 & 0 & 1 \end{pmatrix},\ \rho((23)) = \begin{pmatrix} 1 & 0 & 0 \\ 0 & 0 & 1 \\ 0 & 1 & 0 \end{pmatrix},\ \rho((34)) = \begin{pmatrix} 0 & -1 & 0 \\ -1 & 0 & 0 \\ 0 & 0 & 1 \end{pmatrix}.$$

For our order-3 element we take

$$\sigma = \rho((123)) = \begin{pmatrix} 0 & 0 & 1 \\ 1 & 0 & 0 \\ 0 & 1 & 0 \end{pmatrix}.$$

These matrices permute the vertices of a tetrahedron,

$$w_1 = (-1,+1,+1)\,, \quad w_2 = (+1,-1,+1)\,, \quad w_3 = (+1,+1,-1)\,, \quad w_4 = (-1,-1,-1)\,,$$

according to the natural action of S_4 on a set of size 4.

As stated above, our ansatz (23) and the conditions (27) implied by (25) already reduce our search a great deal. Specifically, we need to choose a unit vector y in the plane R_σ rotated by σ, giving us a free parameter θ, as in Section 3. This plane is perpendicular to w_4, so we have

$$u = y + aw_4 \quad \text{and} \quad v = z + bw_4\,. \tag{28}$$

Moreover, $\langle u_\perp\,|\,v_\perp\rangle = 0$, so either $a = 0$ or $b = 0$, giving two families with two free parameters. Here we derive additional constraints using the Fourier technique we used for $n = 2$. These will suffice to reduce u and v to a finite number of possibilities, which indeed give valid matrix multiplication algorithms.

Let us denote the irreducible representations of S_4 by $\mathbf{1}$ (trivial), π (sign), β for the 2-dimensional irreducible representation, ρ for the standard representation, and $\rho' = \rho \otimes \pi$ for the "anti-standard" representation. Using a little character theory, we find that $\rho \otimes \rho^* \cong \rho \otimes \rho \cong \mathbf{1} \oplus \beta \oplus \rho \oplus \rho'$. Specifically, $\mathbf{1}, \beta, \rho$, and ρ' consist respectively of the scalars, the diagonal matrices with trace zero, the symmetric matrices that are zero on the diagonal, and the antisymmetric matrices. With some additional choices, this decomposition into irreps yields the following orthogonal basis of $\rho \otimes \rho^*$:

$$\begin{aligned}
\mathbf{1} &= \left\langle \mathbf{1} = \begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix} \right\rangle, \quad \beta = \left\langle \beta_x = \begin{pmatrix} 1 & 0 & 0 \\ 0 & -1/2 & 0 \\ 0 & 0 & -1/2 \end{pmatrix}, \ \beta_y = \begin{pmatrix} 0 & 0 & 0 \\ 0 & \sqrt{3}/2 & 0 \\ 0 & 0 & -\sqrt{3}/2 \end{pmatrix} \right\rangle \\
\rho &= \left\langle s_1 = \begin{pmatrix} 0 & 0 & 0 \\ 0 & 0 & 1 \\ 0 & 1 & 0 \end{pmatrix}, \ s_2 = \begin{pmatrix} 0 & 0 & 1 \\ 0 & 0 & 0 \\ 1 & 0 & 0 \end{pmatrix}, \ s_3 = \begin{pmatrix} 0 & 1 & 0 \\ 1 & 0 & 0 \\ 0 & 0 & 0 \end{pmatrix} \right\rangle \\
\rho' &= \left\langle a_1 = \begin{pmatrix} 0 & 0 & 0 \\ 0 & 0 & 1 \\ 0 & -1 & 0 \end{pmatrix}, \ a_2 = \begin{pmatrix} 0 & 0 & -1 \\ 0 & 0 & 0 \\ 1 & 0 & 0 \end{pmatrix}, \ a_3 = \begin{pmatrix} 0 & 1 & 0 \\ -1 & 0 & 0 \\ 0 & 0 & 0 \end{pmatrix} \right\rangle
\end{aligned} \tag{29}$$

We chose s_1, s_2, s_3 so that they correspond to the original basis $e_1 = (1,0,0), e_2 = (0,1,0), e_3 = (0,0,1)$ of ρ. Similarly, since $\rho' \cong \rho \otimes \pi$, we chose a_1, a_2, a_3 to be compatible with this mapping in the basis of ρ that we started with.

The Klein four-group $K_4 \cong \mathbb{Z}_2 \times \mathbb{Z}_2$, which appears in the derived series of S_4 and is generated by (12)(34) and (13)(24), consists of the diagonal matrices with an even number of -1s. Since these commute with β_x and β_y and we are applying the conjugation action to $\rho \otimes \rho^*$, we recover the fact that K_4 is the kernel of β, and that β is the standard irrep of the quotient $S_4/K_4 \cong S_3$. We chose β_x and β_y so that they correspond with same basis vectors $(1,0), (0,1)$ of this irrep as in Section 3, and in particular of the copy of $S_3 \subset S_4$ that fixes w_4 and rotates and reflects R_σ.

To derive our additional equations, we will use (23) to compute $\langle \mathrm{MM} \,|\, A \otimes B \otimes C \rangle = \operatorname{tr} ABC$ for certain triples A, B, C of basis elements from (29). Recall that since MM is fixed under the diagonal conjugation action of S_4, it can only have a nonzero projection into triples $\mu \otimes \nu \otimes \lambda$ that contain a copy of the trivial irrep of the diagonal action; thus only basis elements drawn from such triples are informative.

One such triple is $\mathbf{1} \otimes \rho \otimes \rho$, which contains a copy of the trivial irrep spanned by

$$\mathbf{1} \otimes (s_1 \otimes s_1 + s_2 \otimes s_2 + s_3 \otimes s_3) .$$

Computing $\langle \mathrm{MM} \,|\, \mathbf{1} \otimes s_1 \otimes s_1 \rangle = \operatorname{tr} s_1^2$ using (23) and recalling $\langle v \,|\, u \rangle = -1$ from (25) yields

$$2 = 0 + \langle v \,|\, u \rangle \sum_{g \in S_4} \langle v | s_1^g | u \rangle \langle v | s_1^{g\sigma} | u \rangle = - \sum_{g \in S_4} \langle v | s_1^g | u \rangle \langle v | s_1^{g\sigma} | u \rangle ,$$

where we adopt the shorthand $s^g = \rho(g)^\dagger s \rho(g)$ and (abusing notation slightly) $s^{g\sigma} = \sigma^\dagger \rho(g)^\dagger s \rho(g) \sigma$. Note that the 0 comes from $\left\langle \mathbf{1}^{\otimes 3} \,|\, \mathbf{1} \otimes s_1 \otimes s_1 \right\rangle = (\operatorname{tr} \mathbf{1})(\operatorname{tr} s_1)^2$.

Since we chose s_1, s_2, s_3 so that s_i corresponds with e_i in our original basis for ρ, we have $s_1^g = (-1)^{\epsilon_i} s_i$ where the first row of $\rho(g)$ is $(-1)^{\epsilon_i} e_i$. Similarly, $s_1^{g\sigma} = (-1)^{\epsilon_j} s_j$ where the first row of $\rho(g)\sigma$ is $(-1)^{\epsilon_j} e_j$. However, since σ is a permutation matrix with no negative entries, $\epsilon_j = \epsilon_i$ and the signs cancel out. Moreover, multiplying $\rho(g)$ by σ on the right permutes the e_j: that is, $j = (1,2,3)^{-1} i \neq i$. As a result, each distinct pair i, j of indices appears in 8 of the 24 terms in the sum, and we are left with

$$\langle v | s_1 | u \rangle \langle v | s_2 | u \rangle + \langle v | s_2 | u \rangle \langle v | s_3 | u \rangle + \langle v | s_3 | u \rangle \langle v | s_1 | u \rangle = -\frac{1}{4} . \tag{30}$$

We reason similarly in $\mathbf{1} \otimes \rho' \otimes \rho'$, where the trivial irrep is spanned by

$$\mathbf{1} \otimes (a_1 \otimes a_1 + a_2 \otimes a_2 + a_3 \otimes a_3) .$$

Computing $\langle \mathrm{MM} \,|\, \mathbf{1} \otimes a_1 \otimes a_1 \rangle = \operatorname{tr} a_1^2$ using (23) yields

$$-2 = 0 + \langle v \,|\, u \rangle \sum_{g \in S_4} \langle v | a_1^g | u \rangle \langle v | a_1^{g\sigma} | u \rangle = - \sum_{g \in S_4} \langle v | a_1^g | u \rangle \langle v | a_1^{g\sigma} | u \rangle \,.$$

Recall that $\rho'(g) = \pi(g)\rho(g)$ where $\pi(g) = \pm 1$ is the parity. Since (123) is even, g and $g(123)$ have the same parity, so the parity disappears from each term in the sum. Following the same reasoning then gives

$$\langle v | a_1 | u \rangle \langle v | a_2 | u \rangle + \langle v | a_2 | u \rangle \langle v | a_3 | u \rangle + \langle v | a_3 | u \rangle \langle v | a_1 | u \rangle = \frac{1}{4} \,. \tag{31}$$

We will derive one more constraint along these lines. Analogous to (7), $\beta \otimes \beta \otimes \beta$ contains a copy of the trivial irrep spanned by

$$\beta_x \otimes \beta_x \otimes \beta_x - \beta_x \otimes \beta_y \otimes \beta_y - \beta_y \otimes \beta_x \otimes \beta_y - \beta_y \otimes \beta_y \otimes \beta_x \,.$$

We compute $\langle \mathrm{MM} \,|\, \beta_x \otimes \beta_x \otimes \beta_x \rangle = \operatorname{tr} \beta_x^3$, giving

$$\frac{3}{4} = 0 + \sum_{g \in S_4} \langle v | \beta_x^g | u \rangle \langle v | \beta_x^{g\sigma} | u \rangle \langle v | \beta_x^{g\sigma^2} | u \rangle \,.$$

Recall that β is the standard irrep of $S_4/K_4 \cong S_3$. Moreover, β_x is stabilized by the transposition $\rho((23))$, so it has an orbit of size 3. Specifically, all its conjugates are either β_x, β_x^σ, or $\beta_x^{\sigma^2}$, and these differ only in which diagonal entry is 1. Each of the 24 terms in the sum contains all three of these, so we have

$$\langle v | \beta_x | u \rangle \langle v | \beta_x^\sigma | u \rangle \langle v | \beta_x^{\sigma^2} | u \rangle = \frac{1}{32} \,. \tag{32}$$

With (30), (31), and (32) in hand, we can further restrict the two-parameter families described by (28) until only valid decompositions of the matrix multiplication tensor remain. First we write a one-parameter family of unit vectors $y \in R_\sigma$,

$$y = y(\theta) = -\sqrt{\frac{2}{3}} \big(\cos\theta, \cos(\theta - 2\pi/3), \cos(\theta - 4\pi/3) \big) \,.$$

in which case

$$z = \frac{2}{3}(\sigma y - y) = -\frac{2}{\sqrt{3}}\, y(\theta + 5\pi/6) \,.$$

Since either $a = 0$ or $b = 0$ is zero in (28), we start with

$$u = y + a w_4 \quad \text{and} \quad v = z \,. \tag{33}$$

Plugging this into (31) yields a constraint on a which is independent of θ,

$$\frac{1}{3} - 2a^2 = \frac{1}{4} \quad \text{so} \quad a = \pm \frac{1}{2\sqrt{6}} \,.$$

If we set $a = -1/(2\sqrt{6})$, then (30) becomes

$$1 - \cos 3\theta - \sqrt{3} \sin 3\theta = 0 \,.$$

This implies that $\theta = (2\pi/3)k$ or $(2\pi/3)(k + 1/3)$ for integer k. The second of these is eliminated by (32), which becomes

$$18 \cos 3\theta + 7\sqrt{3} \sin 3\theta - 8\sqrt{3} \sin 6\theta = 18 \,.$$

Thus θ is an integer multiple of $2\pi/3$. In particular, with $\theta = 0$ we have

$$y = \frac{1}{2\sqrt{6}}(2w_1 - w_2 - w_3)\,,$$

and setting $a = -1/(2\sqrt{6})$ and using $w_4 = -(w_1 + w_2 + w_3)$ gives

$$u = y - \frac{1}{2\sqrt{6}}w_4 = \frac{1}{2}\sqrt{\frac{3}{2}}\,w_1 = \frac{1}{2}\sqrt{\frac{3}{2}}(-1, +1, +1) \quad \text{and} \quad v = z = \frac{1}{\sqrt{6}}(w_2 - w_1) = \sqrt{\frac{2}{3}}(1, -1, 0)\,.$$

This gives an algorithm for $n = 3$ matrix multiplication formed from the rank-1 matrix

$$m = |u\rangle\langle v| = \frac{1}{2}\begin{pmatrix} -1 & +1 & 0 \\ +1 & -1 & 0 \\ +1 & -1 & 0 \end{pmatrix}\,. \tag{34}$$

Taking (33) with $a = 1/(2\sqrt{6})$ gives solutions whenever $\theta = (2\pi/3)(k + 1/2)$, leading to alternate rank-1 matrices which are conjugate to this one.

We also explored the family where $u = y$ and $v = z + bw_4$. In this case, (31) implies that $b = \pm 1/(3\sqrt{2})$. Setting $b = -1/(3\sqrt{2})$ yields solutions at $\theta = (2\pi/3)(k + 3/4)$, and setting $b = 1/(3\sqrt{2})$ yields solutions at $\theta = (2\pi/3)(k + 1/4)$. The resulting rank-1 matrices are conjugates of

$$m' = \frac{1}{2}\begin{pmatrix} -1 & +1 & +1 \\ 0 & 0 & 0 \\ +1 & -1 & -1 \end{pmatrix}\,. \tag{35}$$

These are conjugates of the transpose of the first family of solutions (34). Specifically,

$$m' = \tau m^\dagger \tau$$

where τ is a transposition; in this case, $\tau = \rho((23))$. Indeed, an easy exercise shows that for any m which gives a valid solution to (23), $\tau^{-1}m^\dagger\tau$ does as well, as long as $\tau^{-1}\sigma\tau = \sigma^{-1}$. We leave this to the reader.

4.2 $n = 4$: the standard representation of A_5

For larger n, the pattern continues. Rather than writing out the decomposition into irreps explicitly, we simply summarize the results of our calculations.

For $n = 4$, we let ρ be the standard representation of A_5, which is of order $|A_5| = 60 = n^3 - n$. We construct unit vectors $w_1, \ldots, w_5$ at the corners of a simplex which are permuted by A_5. The two-dimensional subspace R_σ rotated by $\sigma = \rho((123))$ is perpendicular to w_4 and w_5, and in particular to $w_1 + w_2 + w_3 = -(w_4 + w_5)$. In this plane we define

$$y = \sqrt{\frac{2}{15}}\,(2w_1 - w_2 - w_3) \quad \text{and} \quad z = \frac{2}{3}(\sigma y - y) = 2\sqrt{\frac{2}{15}}\,(w_2 - w_1)\,.$$

Using a mixture of numerical search and computational algebra, we found a solution of the form (27), which we do not claim is unique:

$$u = y + \sqrt{\frac{2}{15}}\,(w_1 + w_2 + w_3) = \sqrt{\frac{6}{5}}\,w_1 \quad \text{and} \quad v = z = 2\sqrt{\frac{2}{15}}\,(w_2 - w_1)\,.$$

Unlike the previous case, this representation is not monomial, so it's difficult to find a basis which is both elegant and explicit. One basis, obtained by projecting the 5 unit vectors in $\mathbb{R}^5$ along Fourier basis vectors

with nonzero frequencies, places the corners of the simplex at

$$\begin{aligned}
w_1 &= \frac{1}{\sqrt{2}}\,(1,0,1,0) \\
w_2 &= \frac{1}{4}\left(\frac{\sqrt{5}-1}{\sqrt{2}}, \sqrt{5+\sqrt{5}}, -\frac{\sqrt{5}+1}{\sqrt{2}}, \sqrt{5-\sqrt{5}}\right) \\
w_3 &= \frac{1}{4}\left(-\frac{\sqrt{5}+1}{\sqrt{2}}, \sqrt{5-\sqrt{5}}, \frac{\sqrt{5}-1}{\sqrt{2}}, -\sqrt{5+\sqrt{5}}\right) \\
w_4 &= \frac{1}{4}\left(-\frac{\sqrt{5}+1}{\sqrt{2}}, -\sqrt{5-\sqrt{5}}, \frac{\sqrt{5}-1}{\sqrt{2}}, \sqrt{5+\sqrt{5}}\right) \\
w_5 &= \frac{1}{4}\left(\frac{\sqrt{5}-1}{\sqrt{2}}, -\sqrt{5+\sqrt{5}}, -\frac{\sqrt{5}+1}{\sqrt{2}}, -\sqrt{5-\sqrt{5}}\right) .
\end{aligned}$$

In that case we have

$$u = \sqrt{\frac{3}{5}}\,(1,0,1,0) \quad \text{and} \quad v = \frac{1}{\sqrt{3}}\left(-\frac{\sqrt{5}-1}{2}, \frac{1}{2}\sqrt{\frac{\left(3+\sqrt{5}\right)\left(5-\sqrt{5}\right)}{5}}, -\frac{\sqrt{5}+1}{2}, \sqrt{\frac{5-\sqrt{5}}{10}}\right) ,$$

and

$$m = |u\rangle\langle v| = \frac{1}{\sqrt{5}}\begin{pmatrix} -\frac{\sqrt{5}-1}{2} & \frac{1}{2}\sqrt{\frac{(3+\sqrt{5})(5-\sqrt{5})}{5}} & -\frac{\sqrt{5}+1}{2} & \sqrt{\frac{5-\sqrt{5}}{10}} \\ 0 & 0 & 0 & 0 \\ -\frac{\sqrt{5}-1}{2} & \frac{1}{2}\sqrt{\frac{(3+\sqrt{5})(5-\sqrt{5})}{5}} & -\frac{\sqrt{5}+1}{2} & \sqrt{\frac{5-\sqrt{5}}{10}} \\ 0 & 0 & 0 & 0 \end{pmatrix} .$$

4.3 $n = 5$: an irrep of S_5

For $n = 5$, we again look for a group of size $n^3 - n = 120$. A natural choice is S_5, which has two five-dimensional irreps which differ by the parity. We focus on $\rho = \ydiagram{3,2}$, which is one of the irreducible components of the natural permutation action on subsets $U \subset \{1,2,3,4,5\} = [5]$ with $|U| = 2$. It is spanned by the set of functions $f\colon \binom{[5]}{2} \to \mathbb{R}$ such that, for all $i \in \{1,2,3,4,5\}$, $\sum_{U \ni i} f(U) = 0$.

If we express ρ using the Young orthogonal basis [You32], with basis vectors corresponding to the Young tableaux

$$\begin{array}{|c|c|c|}\hline 1&2&3\\\hline 4&5\\\cline{1-2}\end{array}\,, \quad \begin{array}{|c|c|c|}\hline 1&2&4\\\hline 3&5\\\cline{1-2}\end{array}\,, \quad \begin{array}{|c|c|c|}\hline 1&2&5\\\hline 3&4\\\cline{1-2}\end{array}\,, \quad \begin{array}{|c|c|c|}\hline 1&3&4\\\hline 2&5\\\cline{1-2}\end{array}\,, \quad \begin{array}{|c|c|c|}\hline 1&4&5\\\hline 2&3\\\cline{1-2}\end{array}\,,$$

then we have the order-3 element

$$\sigma = \rho((123)) = \begin{pmatrix} 1 & 0 & 0 & 0 & 0 \\ 0 & -\frac{1}{2} & 0 & \frac{\sqrt{3}}{2} & 0 \\ 0 & 0 & -\frac{1}{2} & 0 & \frac{\sqrt{3}}{2} \\ 0 & -\frac{\sqrt{3}}{2} & 0 & -\frac{1}{2} & 0 \\ 0 & 0 & -\frac{\sqrt{3}}{2} & 0 & -\frac{1}{2} \end{pmatrix} .$$

Note that unlike our previous constructions, σ rotates *two* distinct two-dimensional subspaces rather than one. This gives u and v additional degrees of freedom, meaning that (27) might not apply. However, we can leap directly to a solution inspired by an apparent pattern in the examples above.

As the reader may have noticed, in each case above, we had a set W of vectors $w_1, \ldots, w_k$ which form an orbit under the group, and found a decomposition based on the rank-1 matrix $m = |u\rangle\langle v|$ where u and v are proportional to w_1 and $w_2 - w_1$ respectively. In $n = 2, 3, 4$ this orbit had size $|W| = n + 1$, with the w_i pointing to the corners of a simplex which the group permutes. With $G = S_3$, S_4, and A_5, each w_i had a stabilizer subgroup $H \cong \mathbb{Z}_2$, S_3, and A_4 respectively.

The group S_5 does not have any orbits of size 6 in ρ. However, its subgroup A_5 has an orbit $\{w_1, \ldots, w_6\}$ which is sent to $\{-w_1, \ldots, -w_6\}$ by the odd permutations, so S_5 has an orbit of size 12 of the form $\{\pm w_1, \ldots, \pm w_6\}$. This orbit has a stabilizer of size 20 which is determined up to conjugacy in S_5, and there is a unique conjugacy class of subgroups of S_5 of order 20. Namely, identify $\{1, 2, 3, 4, 5\}$ with $\mathbb{F}_5$ arbitrarily, and let H be the subgroup consisting of $\mathbb{F}_5$-affine functions $x \mapsto ax + b$ where $a \in \mathbb{F}_5^*$ and $b \in \mathbb{F}_5$; note that $|H| = 20$. Then restricting S_5 to H breaks ρ into a 4-dimensional irrep of H and a copy of the sign representation, and each one-dimensional subspace spanned by $\pm w_i$ is fixed up to a sign by H or one of its conjugates. (In particular, each w_i is defined by its stabilizer, in the same way that MM itself is.)

Then if we take

$$w_1 = \frac{1}{\sqrt{5}}\left(1, \sqrt{2}, 0, 0, \sqrt{2}\right) \quad \text{and} \quad w_2 = \sigma w_1 = \frac{1}{\sqrt{10}}\left(\sqrt{2}, -1, \sqrt{3}, -\sqrt{3}, -1\right),$$

we find using computational algebra that

$$u = w_1 \quad \text{and} \quad v = \frac{5}{6}(w_2 - w_1) = \sqrt{\frac{5}{6}}\left(0, -\frac{\sqrt{3}}{2}, \frac{1}{2}, -\frac{1}{2}, -\frac{\sqrt{3}}{2}\right)$$

indeed gives a valid decomposition of the matrix multiplication tensor.

5 Lattices and a construction for general n

In all the constructions we have found, the inner products of the $w_i \in W$ are rational: specifically, $\langle w_i \,|\, w_j \rangle = -1/n$ if we normalize the w_i to have unit length. It follows that W generates a lattice, and we can take u and v to be elements of this lattice up to an overall scale factor. This strongly suggests to us that good constructions for larger n might involve lattices whose symmetry group contains G.

In fact, after rescaling u and v, all our constructions can be written as follows, where $w_1, \ldots, w_{n+1}$ form a frame of unit vectors that point to the corners of a simplex:

$$\mathrm{MM}_n = \mathbf{1}^{\otimes 3} + \left(\frac{n}{n+1}\right)^3 \sum_{i,\,j,\,k \text{ distinct}} |w_i\rangle\langle w_j - w_i| \otimes |w_j\rangle\langle w_k - w_j| \otimes |w_k\rangle\langle w_i - w_k| \,. \tag{36}$$

This sum over all $(n+1)n(n-1) = n^3 - n$ distinct ordered triples occurs whenever the group G acts 3-transitively on the w_i. This is true for S_3 in $n = 2$, S_4 in $n = 3$, A_5 in $n = 4$, and (projectively) S_5 in $n = 6$. (Indeed, the 3-transitive permutation groups have been completely classified, see, e. g., [].) Note that this last action relies on the unusual embedding of S_5 in S_6, which we can obtain by taking a standard embedding (say, the copy $S_5 \subset S_6$ that fixes the 6th element) and applying the outer automorphism of S_6.

There is a simple proof that (36) gives a valid matrix multiplication algorithm for any n. First note that we can sum over all $(n+1)^3$ triples i, j, k, since if any of these are equal the summand is zero. We also have

$$\sum_i |w_i\rangle = 0 \quad \text{and} \quad \frac{n}{n+1}\sum_i |w_i\rangle\langle w_i| = \mathbf{1} \,.$$

Then

$$\left(\frac{n}{n+1}\right)^3 \sum_{i,j,k} |w_i\rangle\langle w_j - w_i| \otimes |w_j\rangle\langle w_k - w_j| \otimes |w_k\rangle\langle w_i - w_k|$$

$$= \left(\frac{n}{n+1}\right)^3 \Bigg[\sum_{i,j,k} |w_i\rangle\langle w_j| \otimes |w_j\rangle\langle w_k| \otimes |w_k\rangle\langle w_i|$$

$$- \sum_{i,j,k} \Big[|w_i\rangle\langle w_i| \otimes |w_j\rangle\langle w_k| \otimes |w_k\rangle\langle w_i| + |w_i\rangle\langle w_j| \otimes |w_j\rangle\langle w_j| \otimes |w_k\rangle\langle w_i| + |w_i\rangle\langle w_j| \otimes |w_j\rangle\langle w_k| \otimes |w_k\rangle\langle w_k| \Big] \tag{37}$$

$$+ \sum_{i,j,k} \Big[|w_i\rangle\langle w_j| \otimes |w_j\rangle\langle w_j| \otimes |w_k\rangle\langle w_k| + |w_i\rangle\langle w_i| \otimes |w_j\rangle\langle w_k| \otimes |w_k\rangle\langle w_k| + |w_i\rangle\langle w_i| \otimes |w_j\rangle\langle w_j| \otimes |w_k\rangle\langle w_i| \Big] \tag{38}$$

$$- \sum_{i,j,k} |w_i\rangle\langle w_i| \otimes |w_j\rangle\langle w_j| \otimes |w_k\rangle\langle w_k| \Bigg]$$

$$= \left(\frac{n}{n+1}\right)^3 \sum_{i,j,k} |w_i\rangle\langle w_j| \otimes |w_j\rangle\langle w_k| \otimes |w_k\rangle\langle w_i| - \left(\frac{n}{n+1} \sum_i |w_i\rangle\langle w_i|\right)^{\otimes 3} \tag{39}$$

$$= \mathrm{MM}_n - \mathbf{1}^{\otimes 3}.$$

The mixed terms in lines (37) and (38) disappear because each product has an index that appears only once, and summing over that index gives zero. To see that the first term in (39) is MM_n, first note that since W is an orbit of an irreducible group action, such as the standard irrep of $G = S_{n+1}$, this term has the G^3 symmetries of matrix multiplication discussed in Section 3. This shows that it is a constant times MM_n, and we can check that this constant is 1 by taking the inner product:

$$\left\langle \mathrm{MM}_n \middle| \left(\frac{n}{n+1}\right)^3 \sum_{i,j,k} |w_i\rangle\langle w_j| \otimes |w_j\rangle\langle w_k| \otimes |w_k\rangle\langle w_i| \right\rangle = \left(\frac{n}{n+1}\right)^3 \sum_{i,j,k} \mathrm{tr}\left(|w_i\rangle \langle w_j|\, |w_j\rangle \langle w_k|\, |w_k\rangle \langle w_i|\right)$$

$$= n^3 = \langle \mathrm{MM}_n | \mathrm{MM}_n \rangle \ .$$

This gives a very compact and transparent proof that Strassen's algorithm is a valid decomposition of MM_2, and shows that for all n there is a decomposition of MM_n of tensor rank $n^3 - n + 1$.

Acknowledgments

Parts of this project were inspired in 2015 by discussions with Jonah Blasiak, Thomas Church, Henry Cohn, and Chris Umans, via a collaboration funded by the AIM SQuaRE program, with an additional visit hosted by the Santa Fe Institute. J.A.G. was funded by an Omidyar Fellowship from the Santa Fe Institute during this work, and C.M. was funded partly by the John Templeton Foundation. C.M. also thanks École Normale Supérieure for providing a visiting position during which some of this work was carried out.

References

[AFLG15] Andris Ambainis, Yuval Filmus, and François Le Gall. Fast matrix multiplication: limitations of the laser method. In *STOC '15: 47th Annual ACM Symposium on Theory of Computing*, pages 585–593, 2015. doi:10.1145/2746539.2746554.

[BCS97] Peter Bürgisser, Michael Clausen, and M. Amin Shokrollahi. *Algebraic complexity theory*, volume 315 of *Grundlehren der Mathematischen Wissenschaften [Fundamental Principles of Mathematical Sciences]*. Springer-Verlag, Berlin, 1997. With the collaboration of Thomas Lickteig. doi:10.1007/978-3-662-03338-8.

[BI11] Peter Bürgisser and Christian Ikenmeyer. Geometric complexity theory and tensor rank. In *STOC '11: 43rd Annual ACM Symposium on Theory of Computing*, pages 509–518. ACM, New York, 2011. Preprint of the full version available as arXiv:1011.1350 [cs.CC]. doi:10.1145/1993636.1993704.

[Bin80] Dario Bini. Relations between exact and approximate bilinear algorithms. Applications. *Calcolo*, 17:87–97, 1980.

[Blä99] Markus Bläser. A $\frac{5}{2}n^2$-lower bound for the rank of $n \times n$-matrix multiplication over arbitrary fields. In *FOCS '99: 40th Annual IEEE Symposium on Foundations of Computer Science*, pages 45–50. IEEE Computer Soc., Los Alamitos, CA, 1999. doi:10.1109/SFFCS.1999.814576.

[Blä03] Markus Bläser. On the complexity of the multiplication of matrices of small formats. *J. Complexity*, 19(1):43–60, 2003. doi:10.1016/S0885-064X(02)00007-9.

[Blä13] Markus Bläser. *Fast Matrix Multiplication*. Number 5 in Graduate Surveys. Theory of Computing Library, 2013. doi:10.4086/toc.gs.2013.005.

[Bur14] Vladimir P. Burichenko. On symmetries of the Strassen algorithm. arXiv:1408.6273 [cs.CC], 2014.

[Bur15] Vladimir P. Burichenko. Symmetries of matrix multiplication algorithms. I. arXiv:1508.01110 [cs.CC], 2015.

[CBH11] Nicolas T. Courtois, Gregory V. Bard, and Daniel Hulme. A new general-purpose method to multiply 3×3 matrices using only 23 multiplications. arXiv:1108.2830 [cs.SC], 2011.

[CILO16] Luca Chiantini, Christian Ikenmeyer, J.M. Landsberg, and Giorgio Ottaviani. The geometry of rank decompositions of matrix multiplication I: 2×2 matrices. arXiv:1610.08364 [cs.CC], 2016.

[CW90] Don Coppersmith and Shmuel Winograd. Matrix multiplication via arithmetic progressions. *J. Symbolic Comput.*, 9(3):251–280, 1990. doi:10.1016/S0747-7171(08)80013-2.

[DS13] A. M. Davie and A. J. Stothers. Improved bound for the complexity of matrix multiplication. *Proc. Roy. Soc. Edinburgh Sect. A*, 143:351–369, 2013. doi:10.1017/S0308210511001648.

[HIL13] Jonathan D. Hauenstein, Christian Ikenmeyer, and J. M. Landsberg. Equations for lower bounds on border rank. *Exp. Math.*, 22(4):372–383, 2013. Preprint available as arXiv:1305.0779 [cs.CC]. doi:10.1080/10586458.2013.825892.

[HK71] J. E. Hopcroft and L. R. Kerr. On minimizing the number of multiplications necessary for matrix multiplication. *SIAM J. Appl. Math.*, 20:30–36, 1971. doi:10.1137/0120004.

[JM86] Rodney W. Johnson and Aileen M. McLoughlin. Noncommutative bilinear algorithms for 3×3 matrix multiplication. *SIAM J. Comput.*, 15(2):595–603, 1986. doi:10.1137/0215043.

[Lad76] Julian D. Laderman. A noncommutative algorithm for multiplying 3×3 matrices using 23 multiplications. *Bull. Amer. Math. Soc.*, 82(1):126–128, 1976. doi:10.1090/S0002-9904-1976-13988-2.

[Lan06] J. M. Landsberg. The border rank of the multiplication of 2×2 matrices is seven. *J. Amer. Math. Soc.*, 19(2):447–459, 2006. doi:10.1090/S0894-0347-05-00506-0.

[Lan14] J. M. Landsberg. New lower bounds for the rank of matrix multiplication. *SIAM J. Comput.*, 43:144–149, 2014. doi:10.1137/120880276.

[LG14] François Le Gall. Powers of tensors and fast matrix multiplication. In *Proceedings of the 39th International Symposium on Symbolic and Algebraic Computation*, ISSAC '14, pages 296–303, New York, NY, USA, 2014. ACM. Full version available as arXiv:1401.7714 [cs.DS]. doi:10.1145/2608628.2608664.

[LM16a] J. M. Landsberg and Mateusz Michałek. A $2n^2 - \log(n) - 1$ lower bound for the border rank of matrix multiplication. arXiv:1608.07486 [cs.CC], 2016.

[LM16b] J. M. Landsberg and Mateusz Michałek. On the geometry of border rank algorithms for matrix multiplication and other tensors with symmetry. arXiv:1601.08229 [math.AG], 2016.

[LR16] J. M. Landsberg and Nicholas Ryder. On the geometry of border rank algorithms for $n \times 2$ by 2×2 matrix multiplication. *Exper. Math.*, 2016. In press. Preprint available as arXiv:1509.08323 [cs.NA]. doi:10.1080/10586458.2016.1162230.

[San09] Piotr Sankowski. Maximum weight bipartite matching in matrix multiplication time. *Theoret. Comput. Sci.*, 410(44):4480–4488, 2009. doi:10.1016/j.tcs.2009.07.028.

[Sch81] A. Schönhage. Partial and total matrix multiplication. *SIAM J. Comput.*, 10(3):434–455, 1981. doi:10.1137/0210032.

[Str69] Volker Strassen. Gaussian elimination is not optimal. *Numer. Math.*, 13:354–356, 1969. doi:10.1007/BF02165411.

[Wil12] Virginia Vassilevska Williams. Multiplying matrices faster than Coppersmith–Winograd. In *STOC '12: 44th Annual ACM Symposium on Theory of Computing*, pages 887–898, New York, NY, USA, 2012. ACM. doi:10.1145/2213977.2214056.

[Win71] S. Winograd. On multiplication of 2×2 matrices. *Linear Algebra and Appl.*, 4:381–388, 1971. doi:10.1016/0024-3795(71)90009-7.

[You32] Alfred Young. On quantitative substitutional analysis. *Proceedings of the London Mathematical Society*, s2–34(1):196–230, 1932. doi:10.1112/plms/s2-34.1.196.

[Zwi02] Uri Zwick. All pairs shortest paths using bridging sets and rectangular matrix multiplication. *J. ACM*, 49(3):289–317 (electronic), 2002. doi:10.1145/567112.567114.